Chapter 2: The Physics of Failure

Figure 2.1 represents a distorted conveyor pulley in overload condition. If this happens to a piece of equipment, the parts fail fast. The equipment will run when you press the start button, but not for long. There is little forgiveness when machines are pushed or distorted beyond their design capability. If you want reliable plant and equipment, parts must stay well within their stress limits. Once the available microstructure is overstressed, there is sure to be failure sooner or later. Plant and equipment are only reliable if their parts' materials of construction are fit and healthy. When loads push a part beyond its capability, it fails, and the machine that it is a part of breaks down.

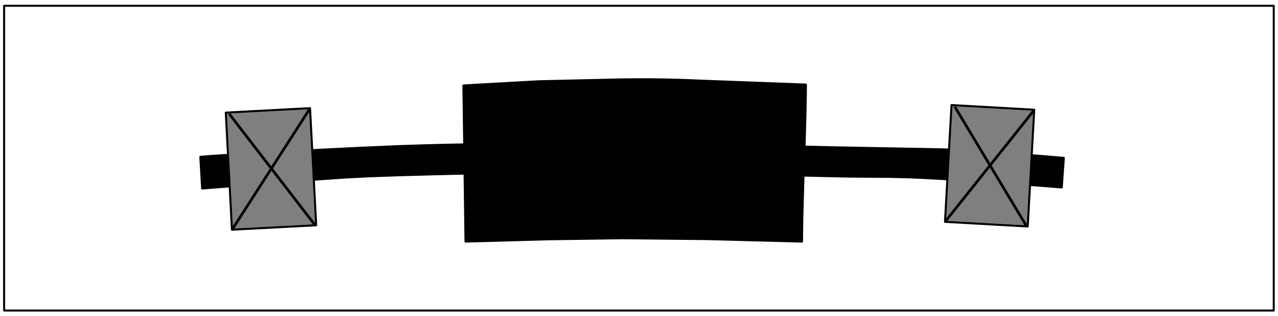

Figure 2.1—Machine Distortion Overloads Parts

There is a retired professor of maintenance and reliability who tells a story in his reliability engineering seminars about the financial outcomes for two organizations with different strategic views on equipment reliability. Some years ago, a maritime operation bought three diesel engines for a new ship. At about the same time, in another part of the world, a railway bought three of the same model diesel engines for a new haulage locomotive. The respective engines went into service on the ship and the locomotive, and no more was thought about either selection.

Several years later, the opportunity arose to compare the costs of using the engines. The ship owners had three times less maintenance cost than the railway. The size of the discrepancy raised interest. An investigation was conducted to find out why there was such a large maintenance cost difference for identical engines in comparable duty. The engines in both services ran for long periods under steady load, with occasional periods of heavier load when the ship ran faster "under-steam" or the locomotive went up inclines. In the end, the difference came down to one factor: the shipping operation had made a strategic decision to de-rate all engines by 10% of nameplate capacity and never run them above 90% of design rating. The railway ran its engines at 100% duty because they had been designed for that duty and so, it was thought, they should be worked at that duty. That single decision reduced the shipping company's maintenance costs by 200%. Such is the size of the financial impact of a seemingly minor difference in the load carried by equipment parts.

The Cause of Machinery and Equipment Parts Failure

A force put on a machine part stresses the part's physical structure. The loads come from the use and operation of plant and equipment under service conditions. Machines break down because their parts' atomic structures can no longer take the imposed forces. Microstructures fail for two reasons: because of deformation, when stress or fatigue causes the atomic bonds to separate, or of degradation, when the atomic bonds are attacked and removed.

When stress is first applied, the atomic matrix reacts by deforming to absorb the load. If the stress is too great, the bonds across the load-carrying section separate and the material breaks—this is *overload*. If a load is excessive but rapidly removed, only a few bonds separate. The unbroken bonds that remain carry the load but with less structure available—this is *fatigue*. The equipment designer's role is to select material for a part with adequate strength for the expected stresses.

Figure 2.2 shows why parts fail from stress and fatigue. It indicates how operating stresses can overload a part's microstructure or weaken it so that it can no longer take the load.

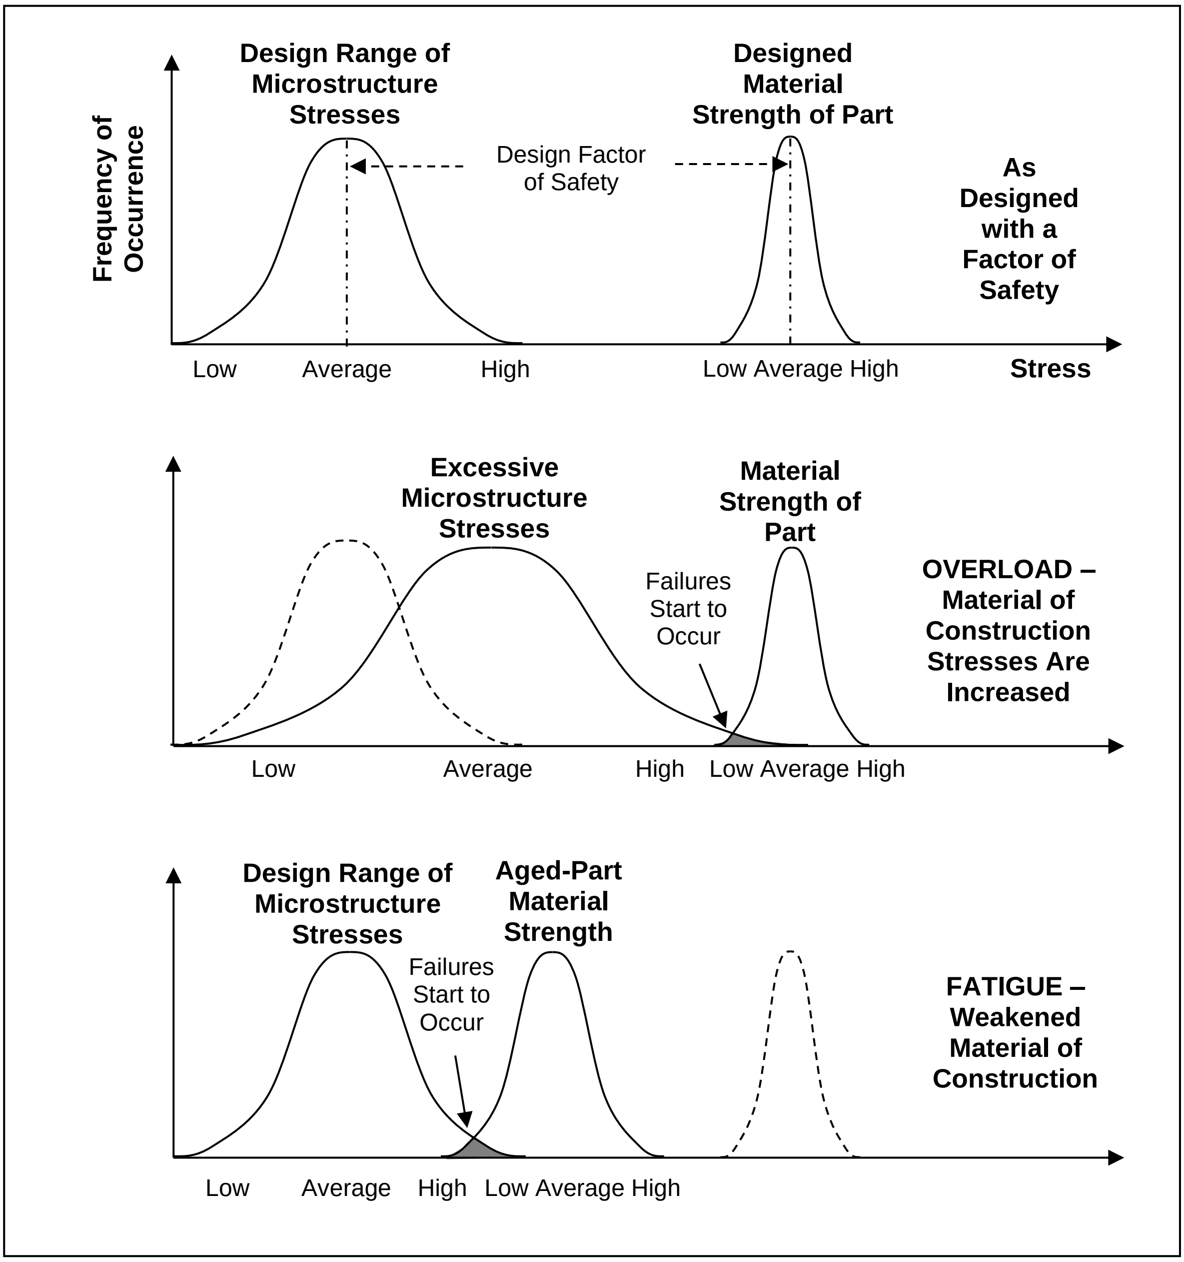

Figure 2.2—Parts Fail When the Stress on Parts Is Greater than the Strength of Parts

The material strength distribution curves on the right-hand side of the figure represent the stress levels at which materials of construction fail. The curves are known as density functions of probability versus stress/strength. They show the natural spread of variation in stress-carrying capacity of identically specified material bought from different suppliers. The load-bearing capacity of a material is dependent on the manufacturer's formulation and how well its chemical and microstructure properties are controlled during formation. This variation is probabilistic, and when you buy material from different suppliers to make parts, there will be a range of stress-carrying capabilities for seemingly identical material. The left-hand curves in Figure 2.2 are probability distributions showing the extent of operating stresses that a part is expected to suffer. They vary from negligible when equipment is at rest to maximum under occasional extreme loading. The stress values used in designing machinery parts are those expected when running at the specified operating conditions defined in the service duly scope for the machine.

The top set of curves show the equipment operated and maintained as the designer intends. The strength of material used in a part and the range of expected operational stresses are wide apart. There is no chance that the part will fail, and it can expect a long working life because the highest operating stress is well below the least-strength part's capacity to handle the stress. The gap between the distributions is a factor of safety that the designer gives us to accommodate the unknown and unknowable.

The middle distribution curves represent the situation in which a part's microstructure stresses rise beyond the factor of safety allowance. Some stresses in the part grow so large that they exceed the remaining material strength, and the item overloads and fails. If the designer wrongly choses material with a low-stress capability for a part when the chance of having overloads is great, then at a future time, a stress will arise that exceeds the capacity of the part. The weakest parts will fail early; the strongest will take more stress before they, too, fail.

Equipment failure can be due to aging of parts when accumulated stress or chemical attack over time weaken the materials of construction. This is shown in the bottom curves, where the part's material properties are degraded until it is too weak to carry high loads and it fails.

Figure 2.3 displays how suffering an excessive stress destroys the material of construction. A portion of the material strength is lost with each high-stress incident as atomic bonds and microstructures break. Figure 2.3 also highlights the failure prediction dilemma—the timing and severity of overload incidents is random. Every component failure has its causes, but when a part is exposed to high-risk situations that could cause it to fail, failure becomes a matter of chance. Expecting to get a full working life from parts is impossible if they suffer out-of-design conditions.

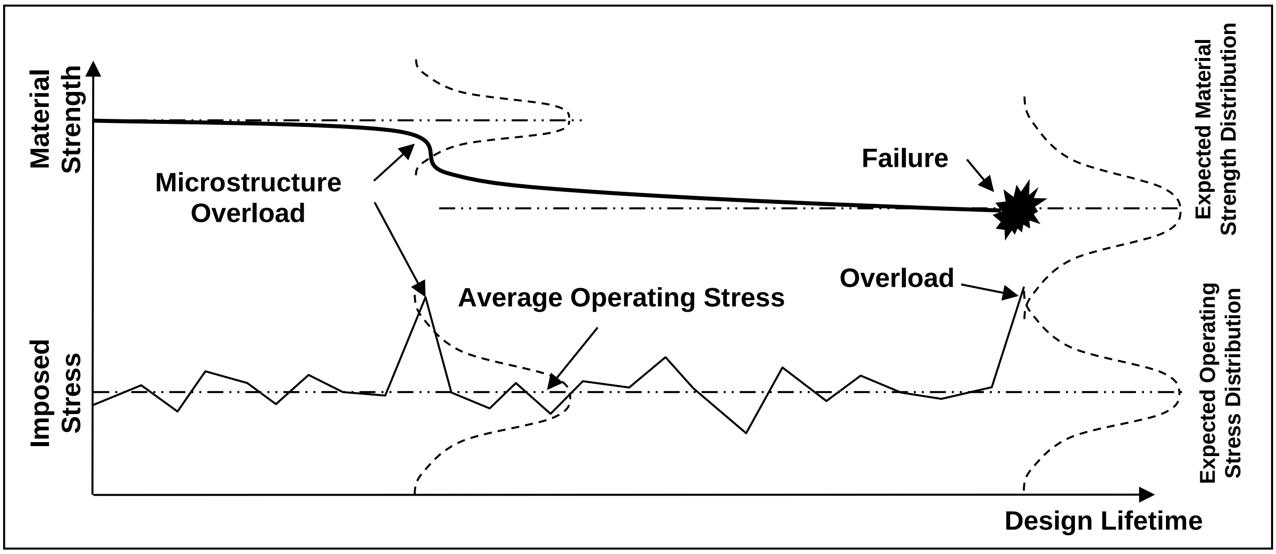

Figure 2.3—Effects of Overload Stresses on the Failure of Parts

Figure 2.4 shows what happens to fatigued parts over the long term. The microstructure gradually weakens, either from the accumulated damage of occasional overload conditions or from the gradual aging by regular fluctuating stresses. Fatigued parts eventually fail because a fateful load occurs one day that destroys the remaining microstructure. These excessive stresses are not necessarily the fault of bad operating practices. They are often attributable to wrong engineering choices or the selection of poor maintenance quality standards that cause stressful situations and increase the probability of failure. If you can prevent distortion of parts so that stress levels on their materials of construction are kept far below the values that break their microstructures, then the parts will not fail, and your machines will be highly reliable—and remain so.

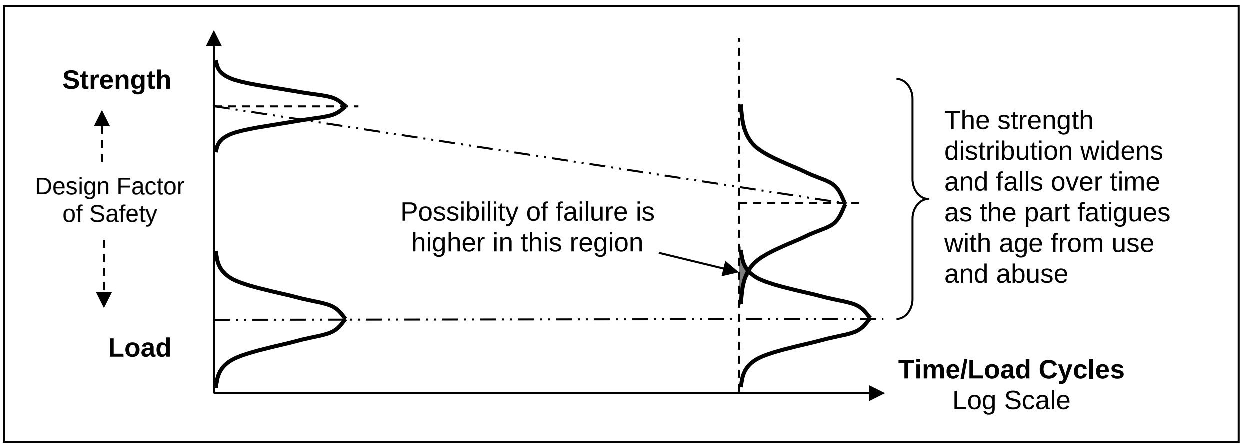

Figure 2.4—Time-Dependent Load and Strength Variation as Stress Damage Accumulates

Parts' microstructures will also fail if chemicals contacting their surface attack the atomic matrix and destroy the matter of their construction. Examples include oxygen in the air degrading rubber; hydrogen ions in water causing steel to corrode through mechanisms such as pitting and crevice corrosion, shown in Figure 2.5; and aggressive chemicals attacking intergranular phases of the microstructure in alloy metals. In these situations, atomic structures fail by degradation. If you prevent degrading environments from touching your equipment parts, then their atomic structure cannot be attacked and failed.

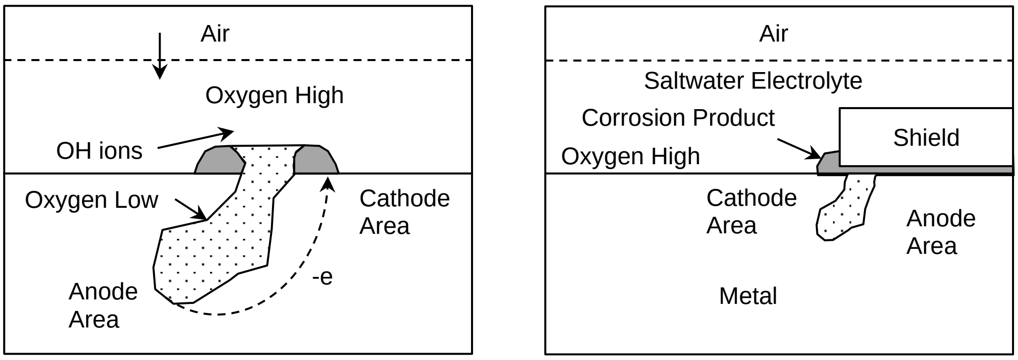

Figure 2.5—Pitting Corrosion and Crevice Corrosion in Seawater

The Physics of Failure

The study of the mechanisms and processes of failure in parts and machines is known as the Physics of Failure (POF). It is the foundation of today's best-practice equipment design methodology used to engineer and build reliable machines. Figure 2.6 shows the POF approach for designing machinery and equipment.[1] It recognizes the influences and effects of the Physics of Failure mechanisms on parts.

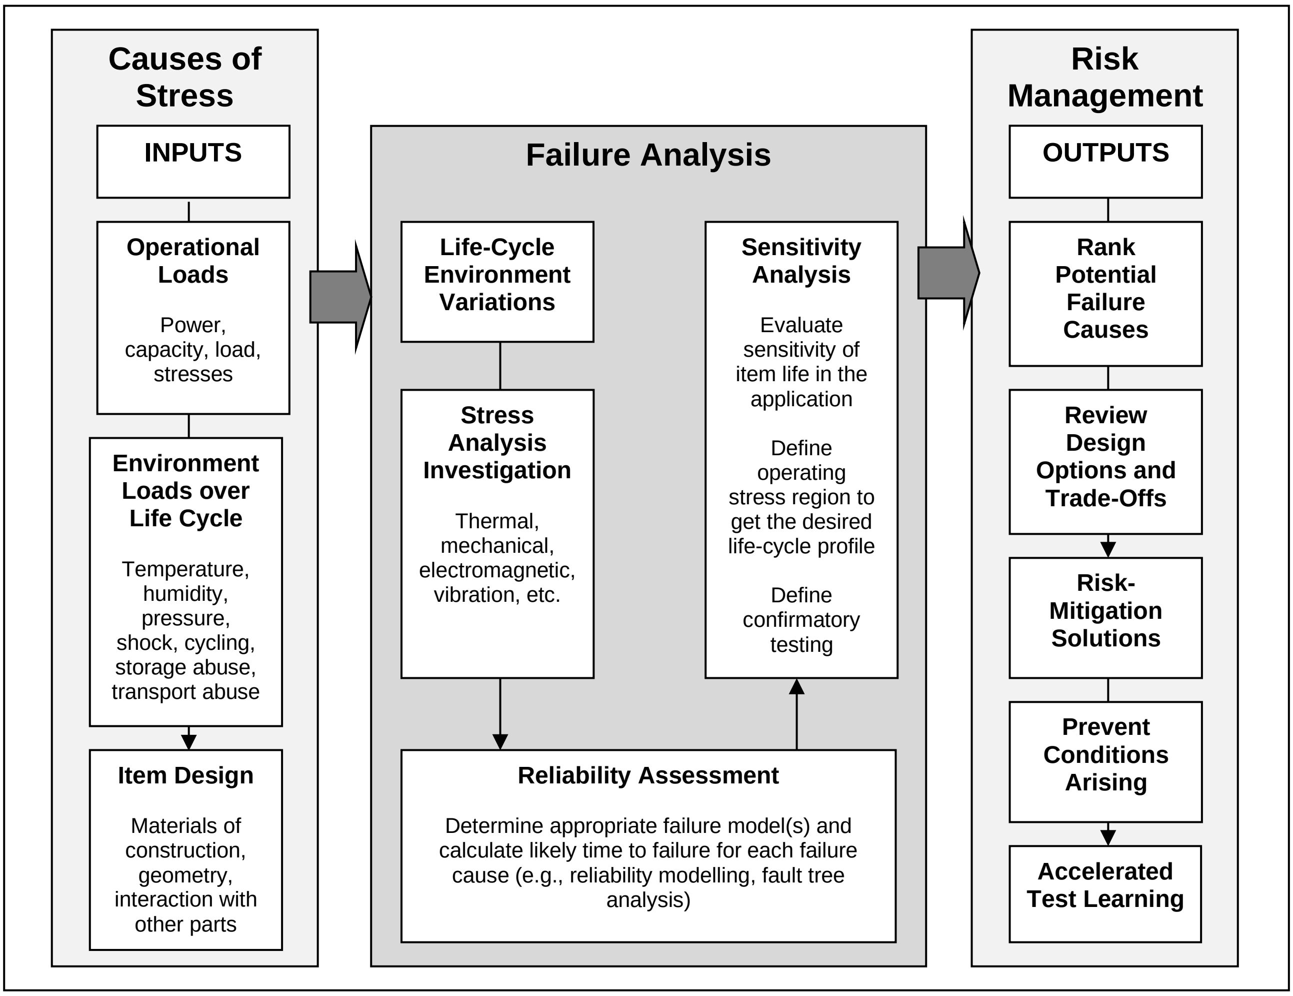

Figure 2.6—Physics of Failure Approach to Reliability Improvement

Equipment components are computer-modeled (or prototypes are laboratory tested) to simulate the performance of the materials of construction in a range of operating situations. The computer-generated model is analyzed for weaknesses. The part is put through various operating conditions, such as overloads, temperature effects, geometry changes, and distortion. The modeling identifies the part's likely behavior in the simulated situations and indicates what loads it can take before failing. The results warn of the design limit and operating envelope for the material of construction. During operation, we must ensure that parts are never loaded and stressed to those levels or allowed to degrade to the point that they cannot take their service loads. It is the role of project engineering, maintenance management, and reliability engineering to ensure that parts do not fail and that machines do not stop during operation.

We know what causes parts to fail and equipment to break down—sudden excessive stress, accumulated stress, or removal of structural material. During plant and equipment design we apply knowledge of the Physics of Failure to select the right materials and configurations to deliver affordable reliability during operating life. The design envelope sets the service limits of a part's working life. To maximize reliability, we first must have parts that can comfortably take extreme service loads. Second, we must ensure that microstructure stresses during operation are kept well within the design envelope. Third, the part's physical structure cannot be attacked by damaging chemicals or elements from the contacting environment. If you want reliable equipment, don't strain parts' microstructures to breaking point or let their materials of construction be harmed.

Limits of Material Strength

Apply force to an object, and it deforms. Its atomic structure is strained. The greater the force that is applied, the more the deformation. In metals, this relationship is known as Hooke's law. Figure 2.7 shows the stress/strain relationship for some common types of engineering metals under constant tensile load. Metals have an elastic region where load and strain are proportional (the straight lines on the graphs). In this region, the metal acts like a spring. Remove the load and the deformation (strain) reduces, and the metal returns to its original shape. If instead the load increases, the strain rises to a point that the microstructure can no longer sustain the load, and the metal yields like plasticine. The yielding can be gradual, as in the left-hand plot, or it can be sudden, as in the right-hand plot.

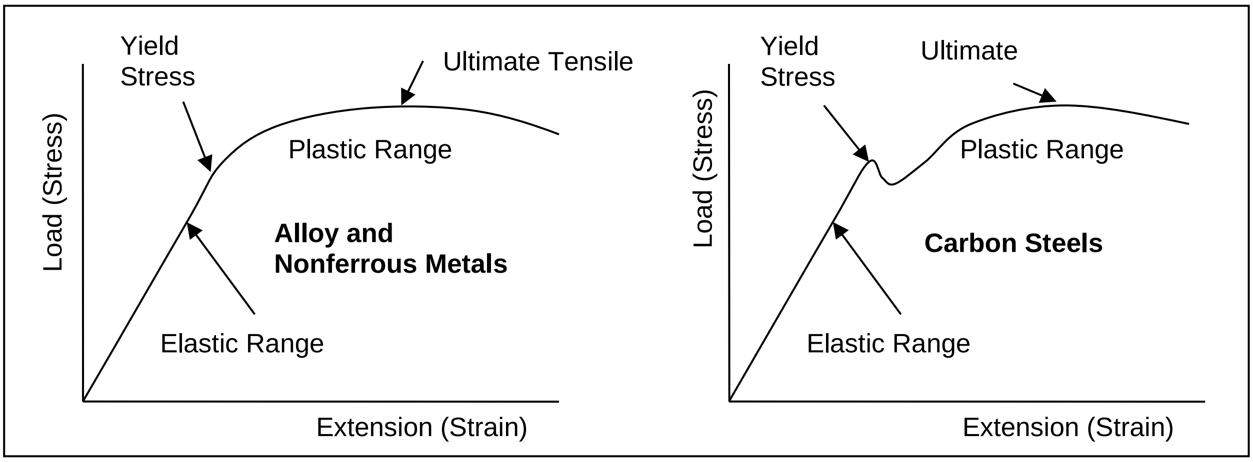

Figure 2.7—When Metal Materials Reach Load Limits, They Deform

Material overstress happens to electronic, electrical, and mechanical parts in equipment that is put under excessive operational loads or suffers environmentally induced stress or forces, such as from vibration, temperature fluctuation, or physical distortion. Parts do not know what causes them stress. They know not the reasons they are overloaded or what caused them to be degraded. Opinions and limited budgets carry no weight with them—they only suffer the facts. Parts react to the stresses and conditions they experience. When the contact region is large, the forces are distributed over a broad surface area and stress is low. As the surface area is reduced, the stresses at the contact points rise. The greatest stress occurs when loads are concentrated on a very small area of microstructure.

After a failure starts in a part, it progresses and grows larger whenever sufficient stress is present. If the stress exceeds the material capability, the part will deform until the microstructures collapse.[2] If a constant stress on a metal part stays well within the elastic region of the curve, the item will not fail. To prevent tensile failure in machine parts, we design for maximum operating stresses that are in the lower portion of the stress/strain line by using factors of safety on yields of 3, 4, and even more for high-risk equipment such as crane hooks and lifting chains. High-tensile bolts are an exception, as they are designed to work in the upper half of the stress/strain line.

Have you ever bent a metal wire or paper clip back and forth until it broke? If you have, you have performed a fatigue stress test. A wire that is bent 90 degrees one way and then 90 degrees the other way does not last long. Each bend produces an overstress, causing damage to accumulate in the microstructure until eventually the wire fatigues and fails.

Fatigue failure is caused by cycling loads. In cyclic stress situations, the microstructure behaves differently than it does under constant tensile or compressive load. When cyclic loading is added into a part's working situation, it markedly reduces the part's life expectancy. This is a confusing phenomenon, as the cycling does not require the load to go from positive to negative—for example, bolt fatigue occurs when the bolt is cycled under positive tensile load. Nor does a cycling load need to exceed factor-of-safety limits for a part to fail. There is also a relationship between a metal's ultimate tensile strength (UTS) and hardness and its ability to handle fatigue loads. The higher the tensile strength and hardness, the more likely it is that the metal will fatigue if it is subject to high fluctuating loads.

There has been a great deal of fatigue load testing done with many materials. Fatigue curves are created in a metallurgical testing laboratory by putting a standard sample bar under cycling tensile load. Eventually the tension combined with cycling stretches the bar to yield by plastic deformation and, ultimately, failure. These tests produce graphs of tensile strength versus number of cycles to failure. They help us understand how much fluctuating load a material can take and still survive. Figure 2.8 is an example of a fatigue curve for wrought (worked) steel, which is used in many industries. Under fluctuating loads just above 90% UTS, the sample lasts 2,000 cycles. Under loads of around 60% UTS, the sample lasts 200,000 cycles before failure—a hundred times longer. But when loads are below half of the UTS, the sample has an indefinite life.

Note that not all metals have a defined fatigue limit as steel does. Nonferrous metals, such as aluminum and brass, have also been fatigue-tested, and their curves are readily available. The fatigue curves for those metals eventually go to failure. Machine parts made of aluminum or brass metals fatigue with use at even low stress levels, and they will need to be replaced well before the part approaches fatigue failure. Fatigue life limits are the reason commercial aircraft frames made of aluminum get retired from service after a specified number of flights. The replacement of parts before failure from operational age and use is known as *preventive maintenance.*

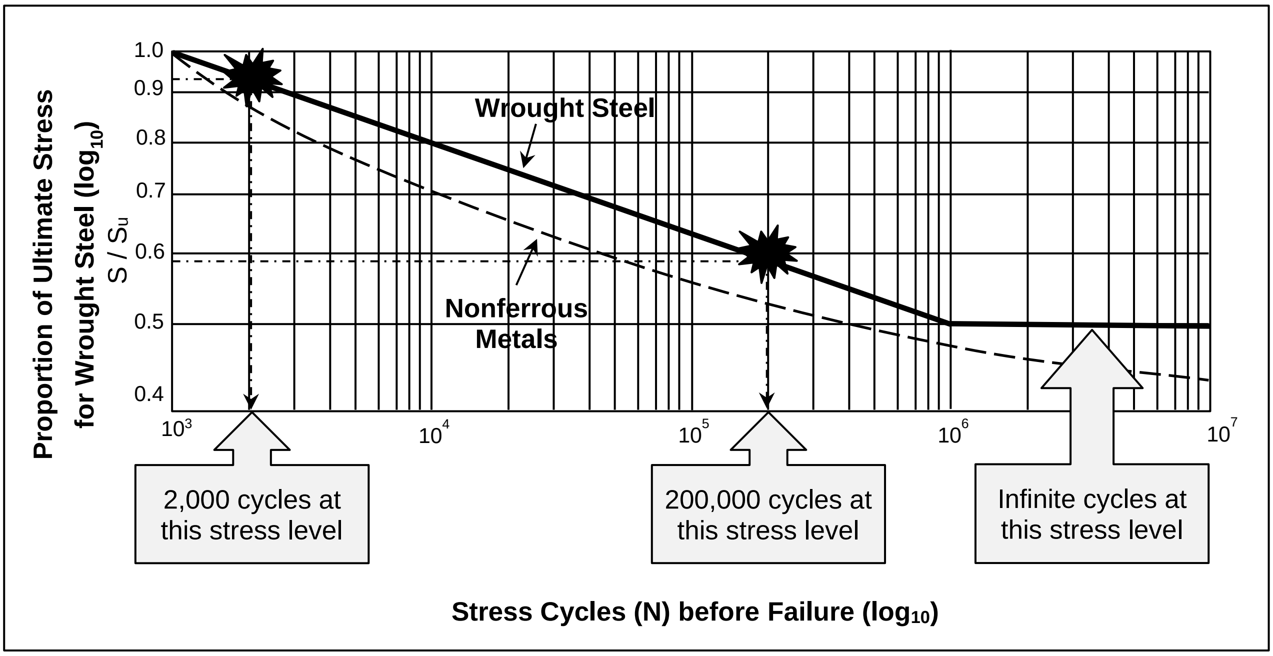

Figure 2.8—Repeated Overstressing Causes Fatigue and Failure

In every case, whether because of extreme burden or cyclic fatigue, excess stress eventually leads to material microstructure failure. If that happens to parts, the equipment will break down.

Engineering Limitations Are a Part of the Design

The loads under which a machine works produce the forces that cause microstructure stresses, which, in turn, case a part to fail. Change the working loads, and the forces change. Change the forces, and the stresses change. We know the engineering formulas of these force/stress relationships. The amount of stress at a particular point of a part's microstructure depends on its shape. To calculate whether there is too much stress on a part, you must know its material properties, its geometry, and the distance between that point and the point where the force is applied. The equations and calculations can get complicated for intricate shapes, but for a round bar, the formulas for the relationships between applied force and the three resultant stresses—axial, torsional, bending—are well documented.

Axial Stress = $F/\pi d^2$

Torsional Stress = $16\ [\ F\ .\ l/\pi d^3\ ]$

Bending Stress = $32\ [\ F\ .\ l/\pi d^3\ ]$,

where F is applied force, d is the bar's diameter, and l is distance to the force from the point being analyzed.

Figure 2.9 shows the change in microstructure stress for a change in applied force on a round bar. All three stresses add together at every point on the bar. With an increase in force, the axial stress rises by an equal proportion. For a torsional load, the microstructure stress rises by 16 times. For a bending load, the stress rises by 32 times. Compounding these stresses is the influence of shape changes in the microstructure, such as the keyways and the shaft fillet, shown in Figure 2.10. Fatigue stresses also increase when parts have sharp contours and when stress raisers, such as notches, holes, and keyways, are present in the part. The contour multiplies the effect of the stress. The stress concentration factor K represents the proportionate rise in the microstructure stress caused by the type of shape change.

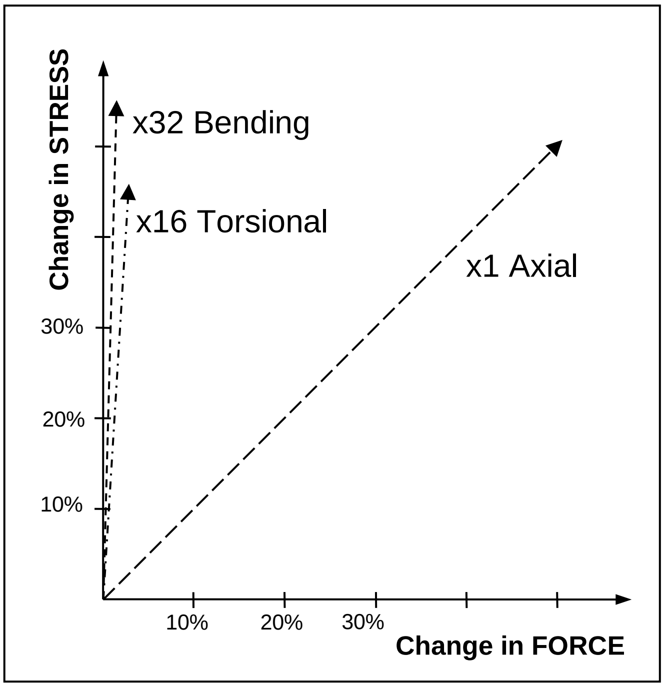

Figure 2.9—Force versus Strain in Round Bars

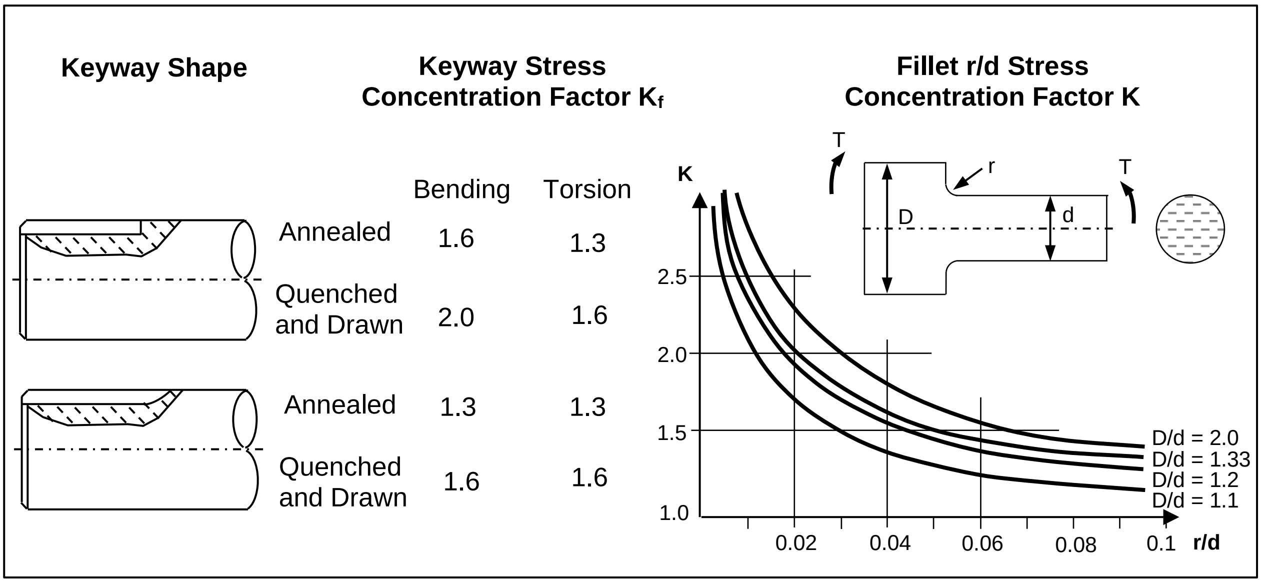

Figure 2.10—Stress Concentration from Keyway Shapes and from Fillets

Because all stresses act concurrently throughout a part, the combination of stresses during high loads may surpass the part's material yield stress at the weakest point and plastically deform the microstructure. If the stress continues past the UTS, a crack will form in the microstructure and dislocations will be created within the material of construction as it separates. Once that happens, things get much worse very fast. The stress needed to propagate an existing failure in the microstructure is significantly less than the stress needed to start the failure. The atomic crack front of a dislocation is razor sharp (imagine you're on the cutting edge of a sharp knife slicing through cheese). Any load that is applied at a crack stress concentration point is multiplied by orders of magnitude.[3] When you change the load conditions on parts by a few percentage points, you can alter stress levels by thousands of percent. Once there are cracks in the microstructure, the resulting failure can be instant. Parts with stress concentration fractures within the material of construction can break from fatigue even under normal operating loads.

Production Limitations Are a Part of the Design

The metallurgical and engineering limits inherent in a machine always dominate production requirements. Machines create fatigue situations for their own parts. As a machine operates, its components are put under tensile and cyclic load, and stress is created in the parts' microstructure. Figure 2.11 is the previous "stress versus cycles to fatigue" failure plot for wrought steel. The dotted lines from the y-axis to the S-N Curve are spaced 10% apart. Where they cross the fatigue curve, the dotted lines drop to the x-axis. At 80% UTS, a sample piece is expected to last about 10,000 cycles before failing. At 70% UTS, it should last about 40,000 cycles. At 60% UTS, its life is about 200,000 cycles before failure. Each 10% drop in stress increases operating fatigue life by four or five times. In fatigue situations, every change in stress greatly impacts the service lifetime of equipment and the resulting maintenance costs.

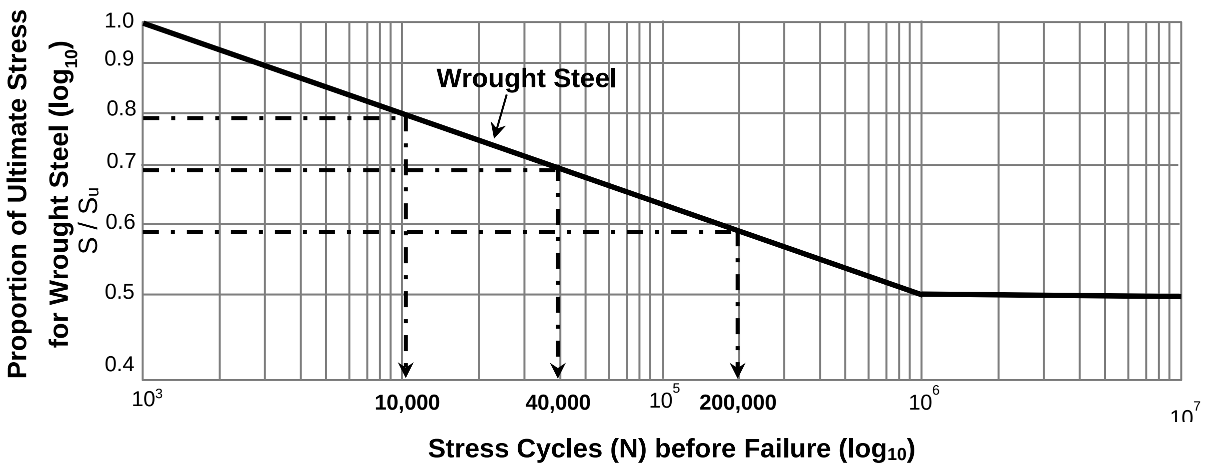

Figure 2.11—Ferrous Metal Fatigue Stress versus Life

For plant and equipment, it is the size of the stresses within parts that dictates whether they retain their integrity. Under load, a part can only survive up to its stress limits; passing those limits leads to its destruction. Somewhere in between disaster and dawdle, you must decide the stress levels at which to run your machines. That decision has surprisingly huge business cost implications because it determines how many times equipment will break down during its life.

Though machine designers limit the effects of stress by using factors of safety in choosing the materials of construction, they also presume that the machine is precisely assembled, that it is installed without extra pre-imposed stresses, and that it is run within its design envelope for its whole life. Situations such as shaft misalignment, soft foot, unbalance, incorrect fits and tolerance, wrong lubricant, insufficient lubricant, water in lubricate, excessive temperatures, and hundreds of other causes of failure are not expected in normal operation. Designers do not engineer equipment to run with parts already heavily preloaded high up the stress/strain line. When machine parts are badly deformed during installation, the microstructure can be so severely stressed that the parts already carry a high percentage of their UTS before the machine even starts operation. Prestressing machine parts because of bad installation will drastically shorten their life—10% additional stress can cost you 10 breakdowns. When machine installers and maintainers put parts and machines into place deformed, they guarantee the early death of the equipment.

During World War II, British airmen referred to ongoing trouble with aircraft, despite aircraft mechanics' best efforts, as "gremlins at work." A gremlin is an imaginary creature that lives inside machines and equipment, breaking parts and causing trouble in systems and devices. We are the gremlins. Our machines and equipment are failed by us in our ignorance. Failure is not an unlucky accident—we murder our machines.[4] They die by being killed in the ways listed below.

- We twist them
- We squeeze them
- We buckle them
- We hit them
- We poison them
- We burn them
- We shake them
- We snap them
- We choke them
- We boil them
- We crush them

Parts do not care why the stress they suffer was imposed on them—they know not that a delivery was late, or that they were overloaded to make up time, or that the repair technician did not have the tools or the skills to do the job right. Parts simply fail when the stress on them becomes excessive, regardless of the cause. Machinery "murderers" are found among equipment manufacturers, among constructors and installers, among equipment users, among maintenance crews, and among business decision makers. People do not intentionally cause failures, but the parts and machines still end up broken.

All machines have production limitations. Run them with parts distorted by overloads or deformed out of shape during installation, and they will soon fail. Overloading a plant to try to make up for lost production time or to fulfill late orders may seem like a heroic management choice, but really it is a destructive asset management and business risk decision. Managers who approve overload levels of operation cause the early death of their machinery. To maximize plant uptime and throughput, run equipment at production rates at which the working parts' microstructure is always below design stresses. Install your machines so they are deformation free. Make sure the surfaces of parts under load have no stress-raising damage. You will be a wise asset manager if you never run your plant harder than 90% of its rated load.

Equipment Reliability Cliffs

Equipment and machinery reliability grows as machines are made more accurately, installed precisely, and components are kept in better condition and health. Once you reach superb operating and maintenance precision with accurate work quality control, your reliability growth will bound forward by years at a time. Conversely, when assembly accuracy is poor or parts' environments are degraded, you will lose reliability at a rate of years at a time. In fact, the rate at which you lose reliability when component health and strength is lost is so steep that it is as if you'd flung your machines over a cliff. Many industrial companies blindly commit industrial suicide daily by leaping off "reliability cliffs." Reliability cliffs show up in graphs and plots of equipment service life as rapid declines into reliability disaster.

Oil Particulate Contamination Reliability Cliff

The table on the left side of Figure 2.12 is famous in the world of equipment maintenance and reliability.[5] It shows the results found by the British Hydromechanics Research Association in the 1990's from a three-year controlled field study of 117 hydraulic machines (injection molding, machine tools, material handling, mobile equipment [earth moving, etc.], marine hydraulics, and test stands) conducted with the purpose of correlating hydraulic fluid cleanliness with machine breakdown frequency. Using a life factor equal to 1 achieved with the ISO 4406 18/15 level of particle contamination often recommended by equipment manufacturers, you can see how quickly operating life falls once hydraulic oil contamination rises above that value. Oil that is dirty with solid particulates rapidly fails hydraulic systems. When you put the information in the table onto the graph on the right-hand side, you see a reliability cliff.

Ave, ISO Code	Ave, Hrs between Breakdowns	Relative Life Factor
24/21	200	0.19
23/20	250	0.24
22/19	325	0.31
21/18	430	0.41
20/17	600	0.57
19/16	800	0.76
18/15	1,050	1
17/14	1,400	1.33
16/13	1,900	1.81
15/12	2,600	2.48
14/11	3,800	3.62
13/10	5,000	4.76
12/9	6,500	6.19
11/8	9,000	8.57
10/7	20,000	19.05

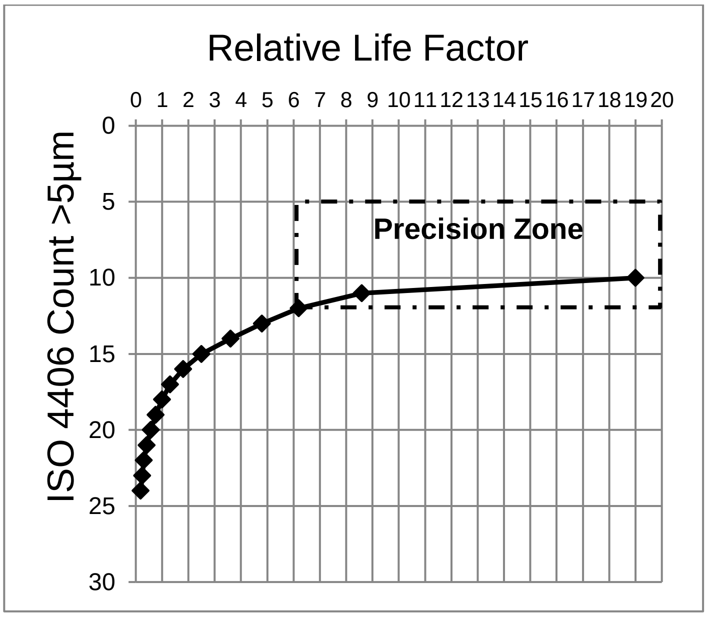

Figure 2.12—Hydraulic Oil Contamination Reliability Cliff

<u>Metal Fatigue Reliability Cliff</u>

The graph on the left side of Figure 2.13 is a laboratory fatigue curve for a wrought steel. The cycles to failure, N, are charted against applied stress, S, producing an S-N fatigue plot. When the data are graphed on the linear scales of the right-hand chart, another reliability cliff appears.

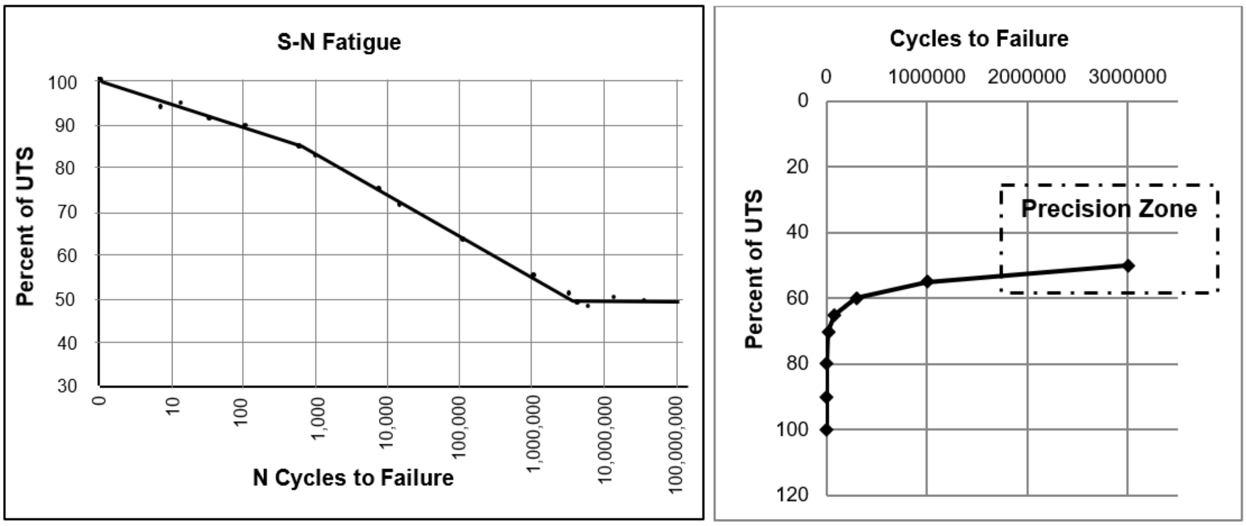

Figure 2.13—Ferrous Metal Fatigue Reliability Cliff]

Roller Bearing Internal Clearance Reliability Cliff

The curve depicted in Figure 2.14 is from laboratory research. It shows the ratio of 6310 ball-bearing service life compared with design life versus bearing internal clearance (25.4 microns, or μm = 0.001 inch). It presents the effect of altering bearing internal clearance on bearing operating lifespan. You can easily see a reliability cliff to the left and the right.

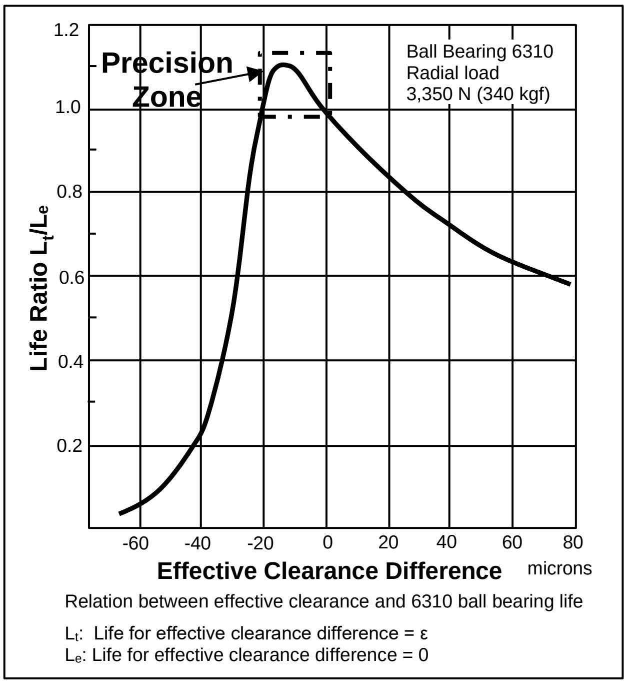

Figure 2.14—Bearing Life versus Bearing Clearance Reliability Cliff

For a 6310 ball-bearing, a small preload improves service lifetime above designed life, but once that point is passed, or if it is not achieved, service lifetime falls off a reliability cliff into rapid bearing failure.

Figure 2.15 derives from the 6310 ball-bearing life curve and shows how much change will destroy each 10% of bearing life. Starting at the full life preload of –20 microns, the change in clearance for each 10% loss of life is gauged and graphed. Another reliability cliff—an abyss, actually—appears. Every 2-micron preload error away from the ideal preload will destroy 10% of the bearing’s service life. Once the bearing is already horribly preloaded, it takes another 5-micron preload to lose 10% more service life. Two microns is about the size of a germ.

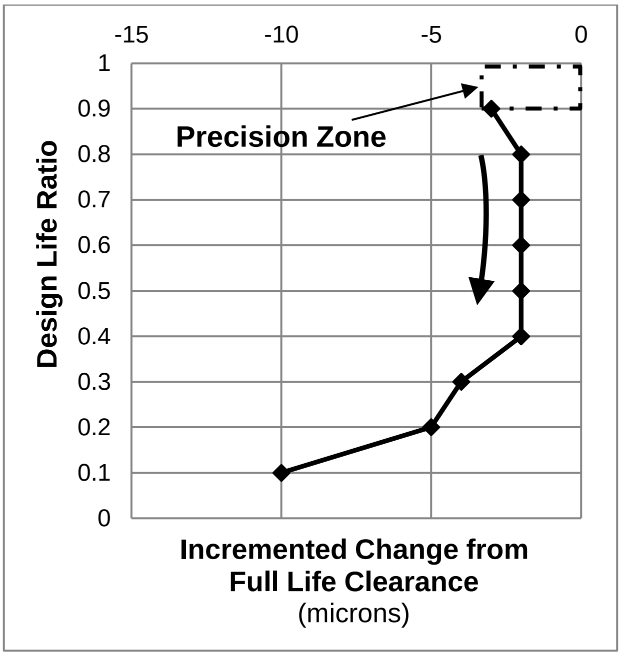

Figure 2.15—Bearing Life versus Change of Bearing Clearance Reliability Cliff

Getting roller bearing clearance right is vital in machinery assembly. The machine designer specifies suitable manufacturing sizes and tolerances for optimum new assembly fits. Keeping internal clearances right—that is, within 10 microns of perfect—during operation and rebuild is enormously hard. You must prevent component distortion, prevent temperature rise (and fall), and minimize fatigue stress cycling: it means you must control all operating conditions masterly. It is hard to do, but it is vital for high roller bearing reliability if you want zero-breakdown machinery.

Shaft Misalignment Reliability Cliff

The left-side graph in Figure 2.16 is taken from the *Shaft Alignment Handbook* by John Piotrowski.[6] It is a plot of the typical operating lifespan of rotating machinery subjected to various amounts of shaft misalignment. The data are based on many case studies in which shaft misalignment was found to be the root cause of rotating equipment breakdown. Graph the bottom curve of the failure zone in the chart, and the reliability cliff on the right reveals itself.

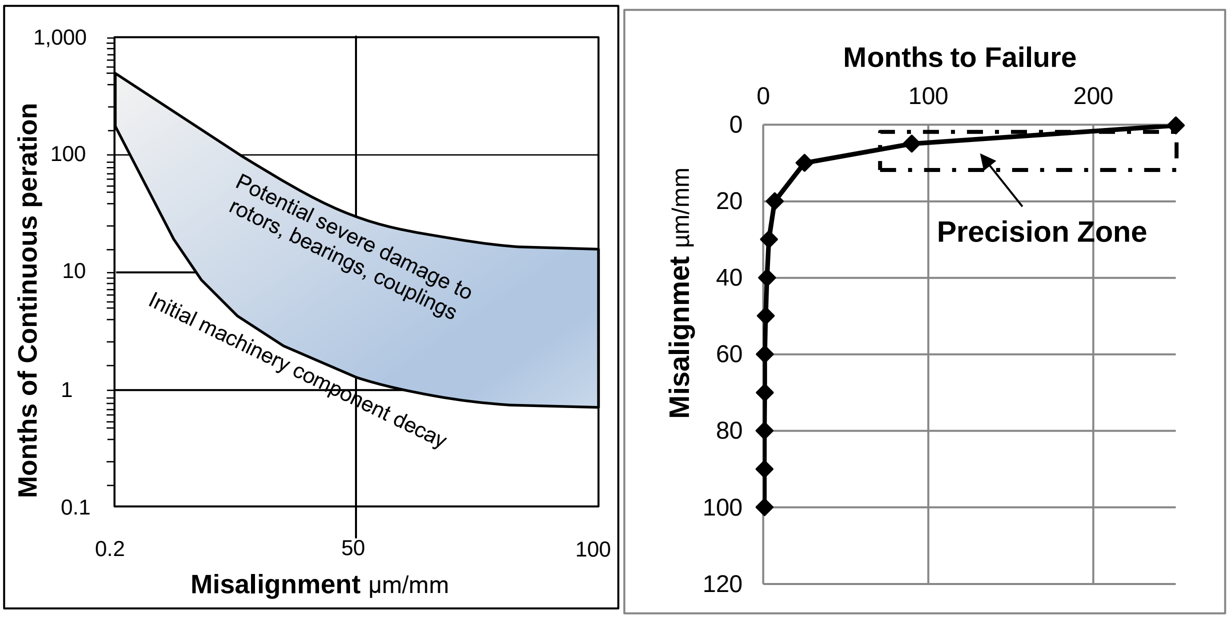

Figure 2.16—Shaft Misalignment versus Machine Life Reliability Cliff

Equipment and Parts Degradation Cycle

Overstressed parts get damaged and then fail. Every overload steals a little of a part's life. The degradation cycle depicted in Figure 2.17 shows the influence of a part's failure sequence on its performance. All parts go through the recognizable stages of degradation. First the component works as expected, but then abnormal situations arise, and it gradually starts to fail.

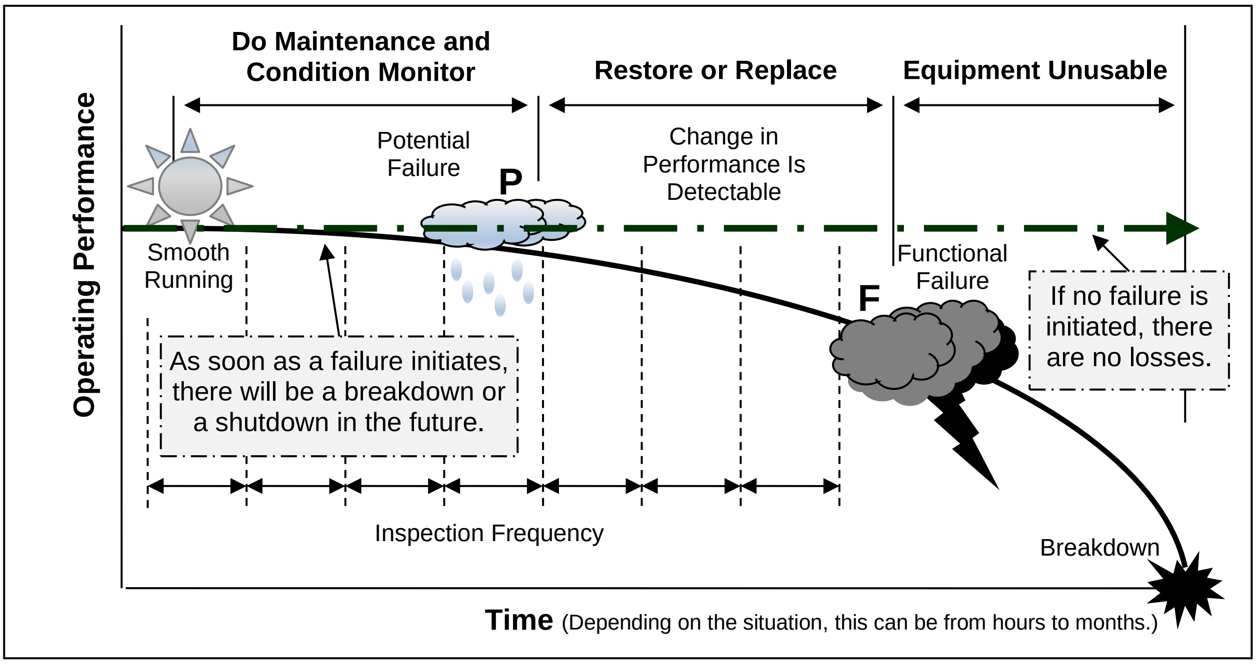

Figure 2.17—Component Failure Degradation Sequence

Some parts fail without exhibiting warning signs of a coming disaster. They show no evidence of degradation—there is just a sudden, catastrophic failure. In such cases, all we see is the sudden death of the part. This commonly happens to electronic parts. It is worth noting that almost all failures, even to electrical and electronic parts, are ultimately mechanical (involving relative movement), contaminant, or over-temperature related.[7] For the most part, we can prevent those situations.

The point at which degradation is first possible to detect is the potential failure point,[8] P. After this point, failure will eventually happen; how soon breakdown occurs depends on the size of the stresses imposed on the failing part. The point at which degradation has progressed beyond salvage and the equipment performance is critically affected is the functional failure point, F.

The degradation cycle is the basis of predictive maintenance strategy and condition monitoring. The degradation curve explains why and when to use condition monitoring on plant and equipment. Knowing that equipment parts show evidence of developing failure it is sensible to inspect them at regular intervals for signs of approaching failure. Once you select an appropriate technology to detect and measure degradation, a part's condition can be trended, and the impending failure monitored until it is time to make a repair.

You set the inspection frequency to detect the onset of a problem so there is time to address the failure before it happens. Condition monitoring can be as simple as regular "feel and listen" observations of parts and equipment performance by the operator or as complicated as using continuous on-board monitoring with sensors and instrumentation feeding computer-controlled diagnostic and prognostic software to judge the likely remaining life. When repair work takes a long time to plan, organize, and schedule, that justifies using more complex technology and giving yourself a long P-F interval to prepare for the restoration or rectification. If you choose to use low-tech condition monitoring solutions with short P-F windows, there will be little time to the breakdown after a failure is detected. Such a predictive maintenance strategy necessitates having quick access to spare parts or having redundant equipment to take over when the duty machine or plant fails.

During operation, every company using physical assets faces the dilemma of probabilistic parts failure, shown in Figure 2.18. The figure presents a series of degradation curves for selected working parts in a centrifugal pump-set. Each part has varying P and F points. The longevity to the P point and the drop rate of the P-F curve for each part depends on the stresses on the materials of construction and the physical degradations suffered by the part. When you use plant and equipment, you live in a situation in which any part can fail at any time if its operating stresses exceed the microstructure's load-carrying capacity. Our machines and equipment typically contain dozens, and often hundreds, of parts under stress, and each of them presents an opportunity to have a breakdown. Every physical asset lives in probabilistic uncertainty because of the multiple opportunities for failure caused by all its degrading and deformed parts. The best protection against failure is to proactively keep every critical part in your plant and equipment fit, healthy, and well for its lifetime.

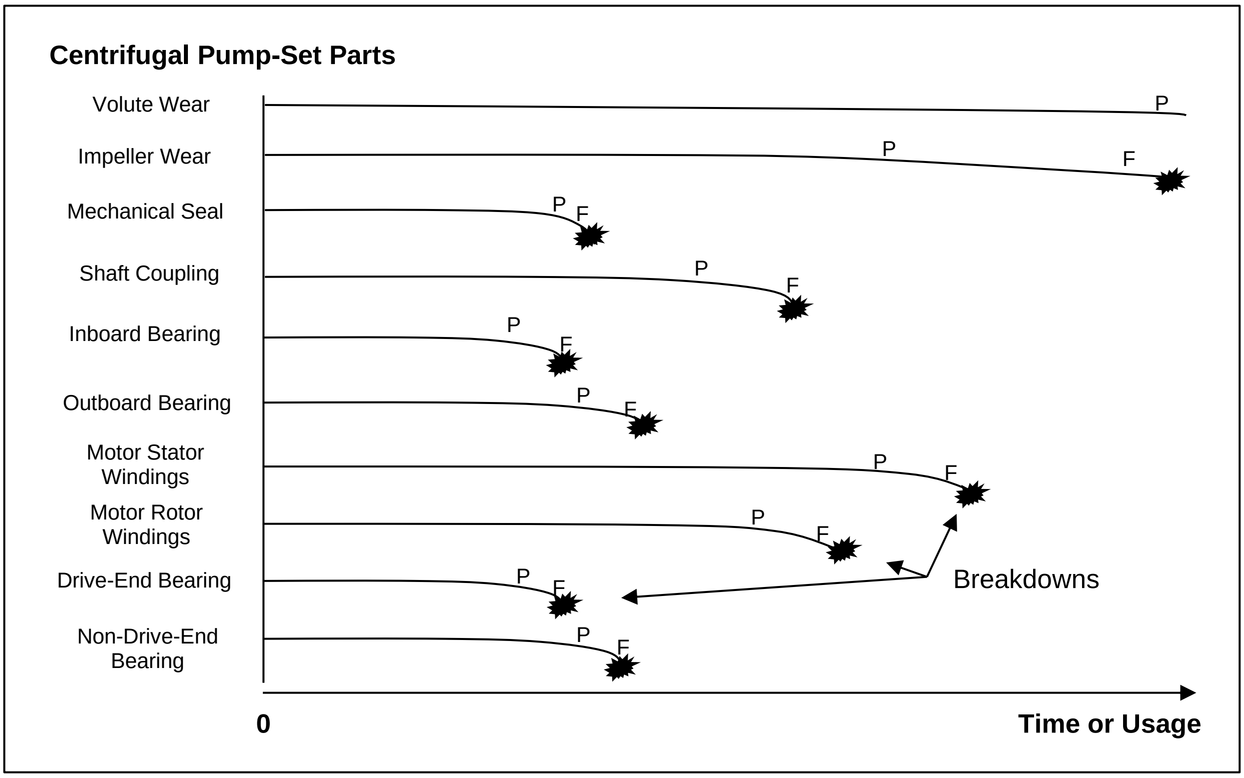

Figure 2.18—Each Part within Equipment Degrades on Its Own Degradation Curve

Reach the Standards That Cause Reliability

Reliability success is often said to be a journey, but that is the wrong metaphor—reliability success is a climb! World-class machinery reliability is only found at the peaks of machinery precision zones. It cannot survive anywhere else: long-lived, trouble-free reliability is a mountainous climb to the place where lasting precision quality lives.

Reliability is a step change. Either you already have reliability, or you are doing one of the following: (1) constantly correcting problems or (2) fixing broken things. There is only one way to make machines reliable: their parts must be even more reliable. The health of your parts produces the reliability of your machines. High machine reliability requires asset life-cycle processes and supply chains that deliver dependable precision quality to working parts. The path away from reliability cliffs starts by setting health quality targets for parts that, when achieved, will result in reliability.

Plant and equipment failure factors such as vibration, fits and tolerances, deformation, unbalance, misalignment, lubricant condition, and fastener tension need to be quality controlled to the values that produce highly reliable machines. Let one go out of control, and world-class reliability is unachievable. Step out of the precision zone with any factor, and you fall over its "reliability cliff" and your machines die. Reliability requires developing design, engineering, manufacture, warehousing, installation, operating, and maintenance procedures with quality standards to prevent deformation and degradation and then training managers, engineers, operators, and maintainers to follow them with great certainty.

Life is wonderful at the top of the reliability cliffs. The view is magnificent. The air is clear and fresh. You get the sunshine from dawn to dusk. You have time and money to enjoy yourself and pursue the happiness of excellence. But take one step too far in the wrong direction, and success ends at the edge of the cliff.

FOOTNOTES

1. Michael Pecht, "Why the Traditional Reliability Prediction Models Do Not Work—Is There an Alternative?" Center for Advanced Life Cycle Engineering, Electronic Product and Systems Center, University of Maryland.

2. J. E. Gordon, *The New Science of Strong Materials, or, Why You Don't Fall through the Floor,* 2nd ed. (New York: Penguin 1976).

3. Robert C. Juvinall, *Engineering Considerations of Stress, Strain, and Strength* (New York: McGraw-Hill, 1967).

4. Rod Bennett, "Machines Don't Die, They're Murdered" (National Condition Monitoring Forum Australia, August 2006).

5. Accessed at http://www.oiltransfer.com/files/british.pdf, June 26, 2015.

6. John Piotrowski, *Shaft Alignment Handbook*, 3rd ed. (Boca Raton, FL: CRC Press, 2007).

7. First heard from retired professor David Sherwin in his three-day "Introduction to Reliability Engineering" course, Perth, Western Australia, June 2007.

8. John Moubray, *Reliability-Centred Maintenance* (Oxford, UK: Butterworth-Heinemann, 1991).